

SYSTEM AND METHOD FOR PROCESSING A REQUEST
USING MULTIPLE DATABASE UNITS

TECHNICAL FIELD OF THE INVENTION

This invention relates in general to the field of
data management, and more particularly to a system and
method for processing a request using multiple database
5 units.

BACKGROUND OF THE INVENTION

The field of data storage has become increasingly important in today's society. Data storage generally involves a database representing a collection of information organized in a way that allows someone or something to access desired pieces of information. One aspect of data storage relates to data management, which generally relates to organizing and accessing information stored in a given database. Data management architectures generally facilitate the retrieval of information or records stored in any one of a number of databases. Additionally, data management architectures may operate to store, to modify, and to update, information in a given database. The operation of such data management architectures continues to grow more challenging as databases grow in size, i.e. the amount of information being stored or maintained in the database.

In order to enable access to data by users or systems at different locations, and to further ensure that data is not lost due to any number of reasons, such as accidents, catastrophes, or operational failures for example, multiple copies of data may need to be effectively managed, maintained, and kept current in certain circumstances.

SUMMARY OF THE INVENTION

From the foregoing, it may be appreciated by those skilled in the art that a need has arisen for an improved data management system for processing a request using multiple database units. In accordance with one embodiment of the present invention, a system and method for processing a request using multiple database units are provided which substantially eliminate or reduce disadvantages and problems of conventional data management techniques.

According to one embodiment of the present invention, a method for processing a request using multiple database units comprises receiving a request for a record and communicating the request to one or more nodes. One or more of the nodes is operable to provide an interface between one or more associated database units and a network. Additionally, one or more of the nodes is operable to communicate with each other. The method further includes identifying one or more target database units that store the record and accessing the record that is stored in one or more of the target database units. The request is then processed based on the record that is stored in one or more of the target database units such that a response to the request is generated. The response is then returned to the request, the response being based on the record that is stored in one or more of the target database units.

Certain embodiments of the present invention provide a number of technical advantages. For example, according to one embodiment of the present invention, a method for processing a request using multiple database units is disclosed that provides the benefits of redundancy and

failover (i.e. a backup or standby protocol) in data storage. This is due, in part, to the apportionment of information among several database units coupled to a network. Thus, in the scenario where a single database
5 unit is inoperable for whatever reason, this deficiency does not affect data storage or retrieval of records in other database units.

Additionally, according to the teachings of the present invention a method for processing a request using
10 multiple database units is provided that offers the benefit of increased capacity. The increased capacity is a result of the use of several databases implemented to replace the use of a single database that is restricted or otherwise constrained by its architecture or size.

15 The embodiments, which follow, may enjoy some, all, or none of these advantages. Other technical advantages may be readily apparent to one skilled in the art from the following figures, description, and claims.

BRIEF DESCRIPTION OF THE DRAWINGS

To provide a more complete understanding of the present invention and features and advantages thereof, reference is made to the following description, taken in
5 conjunction with the accompanying figures, wherein like reference numerals represent like parts in which:

FIGURE 1 is a block diagram of a data management system illustrating a request that is processed using multiple database units coupled to a network; and

10 FIGURE 2 is a flowchart illustrating a series of steps associated with processing a request with multiple database units that are coupled to a network.

DETAILED DESCRIPTION OF THE INVENTION

FIGURE 1 is a block diagram of a data management system 10 illustrating a request 12 that may be processed using multiple database units 16 coupled to a network 14.
5 Each database unit 16 is coupled to an associated node 18, each of the nodes being operable to communicate with each other in order to access a record stored in any one or several database units 16. The record that is stored in database units 16 may be generally accessed in order
10 to satisfy request 12 with a suitable response. FIGURE 1 also illustrates a terminal 20 that generates request 12 to be processed in network 14.

According to the teachings of the present invention, data management system 10 operates to process request 12
15 by querying all nodes 18 for the location of a record that satisfies request 12. After a location for the desired record is identified, these locations may be accessed via communications between nodes 18 and network 14, such that a response to request 12 may be generated
20 and returned to terminal 20. The distribution of data or records amongst database units 16 provides the ability to store a large amount of data that may be accessed via any one of nodes 18. Additionally, the use of several database units 16, which may communicate using suitable
25 peer-to-peer technology, provides failover (i.e. a backup or standby protocol) and redundancy capabilities to data management system 10 without hindering performance characteristics of an associated database coupled to network 14. Moreover, data management system 10 may
30 provide support in a large distributed database environment, for example, offering replication of data

over any number of platforms, protocols, or architectures.

Request 12 is a request generated by a user of terminal 20 that is seeking to access a record entitled
5 'α.' (may reside or be otherwise stored in any number of database units 16.) Alternatively, request 12 could be any interface, computer, server, device, element, or object (or any resultant generated by these elements) operable to communicate a request for a record or for
10 data to network 14. Request 12 is communicated to network 14 where it may then be passed to any one of nodes 18 in an effort to retrieve record α and to further facilitate a response to be returned to terminal 20. Request 12 may include commands or signals for additional
15 operations, other than retrieval, to be performed on record α. Such operations may include print, modify, or delete, for example, or any other suitable operation (inclusive of any number of structured query language (SQL) commands) according to particular needs. Request
20 12 may be communicated over any appropriate network connection, inclusive of wireless communications, to network 14.

Network 14 is a world wide web network in one embodiment, providing an interface between nodes 18 and
25 terminal 20. Network 14 may communicate with one or more nodes 18 or with one or more database units 16 in any manner, including hub-and-spoke, peer-to-peer, or any other suitable protocol. In particular embodiments, network 14 may include one or more engines or other
30 applications that provide communication services or capabilities to a user of terminal 20. Network 14 may comprise one or more local area networks (LANs),

metropolitan area networks (MANs), wide area networks (WANs), and/or any other appropriate internets, intranets, or wireless architectures where appropriate. Network 14 may be coupled to one or more database units 16, via one or more associated nodes 18.

Database units 16 represent storage mediums that may comprise a collection of information for data storage. The term "data storage" is used to refer to any appropriate data source, representation of data, or other organization of data. The data storage within database units 16 may be hierarchical in nature and/or may provide persistent data storage for data management system 10. For example, database units 16 may be a multi-dimensional database that stores data in a hierarchical and multidimensional format or database units 16 may be a representation of data derived by a server or other appropriate component from data stored in a relational database, in memory, or in any other appropriate location or component. Database units 16 may communicate with each other via network 14, nodes 18, or any suitable link operable to exchange information, data or signals. In addition, database units 16 may be coupled to a mainframe, additional networks, or other communications devices or elements where appropriate.

Database units 16 may comprise redundant information, i.e. information that is included in several of database units 16, or alternatively different versions of a record (based on time-stamping or modifications, for example). This redundancy characteristic is illustrated generally in FIGURE 1 by record α as included within several database units 16. Because database units 16 are also capable of storing records or data that is not

included in peer database units 16, a greater capacity for storing information is provided in data management system 10. Database units 16 may store a single record, a portion of a record, or millions of records according to the teachings of the present invention. Database units 16 may be coupled to any device, interface, or component operable to facilitate communications from network 14. In an alternative embodiment, database units 16 may be included within network 14. The term "target database unit" refers to one or more database units 16 that contain or otherwise store a record queried for by request 12. The target database unit 16 is used in generating an appropriate response to request 12. Database units 16 are each coupled to an associated node 18.

Nodes 18 represent a location for receiving request 12, and may further facilitate the retrieval of request 12 from one or more database units 16. Nodes 18 may be an interface, a computer or any other object or element operable to facilitate a communication between network 14 and database units 16. Alternatively, nodes 18 may be coupled to network 14 using one or more local area networks (LANs), metropolitan area networks (MANs), wide area networks (WANs), a global computer network such as the Internet, and/or any other appropriate wireline, wireless, or other links. Additionally, one or more nodes 18 may comprise an application program interface (API) operable to provide a set of routines, operations, protocols, or tools for facilitating a communication between database units 16 and network 14. An API may also provide the ability for various formats, types, or variations of request 12 to be received and processed by

network 14. In one embodiment, nodes 18 are included within associated database units 16. Alternatively, nodes 18 may be included in network 14, facilitating communications between an associated database unit 16 and
5 network 14.

Nodes 18 may operate to perform or otherwise facilitate some operation associated with request 12 such as print, copy, delete, or modify, for example, or any other suitable operation associated with request 12
10 according to particular needs. In the case where operations are performed on request 12, potentially at nodes 18 or alternatively within database units 16, one or more nodes 18 may broadcast the operations to be performed to devices or elements within data management
15 system 10, or alternatively may retrieve the appropriate record (for example record α) and perform the requisite operation on the record before returning it to a suitable location. In the scenario in which one or more nodes 18 perform operations to a given record, an updated copy of
20 the record may be optionally returned to all locations that include the record, or only some locations (that may or may not have included the record) where appropriate and according to particular needs.

There are a variety of ways (as described below) in
25 which one or more nodes 18 may communicate messages or appropriate signals to other nodes 18 or database units 16 such that a targeted record is retrieved and used to satisfy request 12. For example, request 12 may be ported from one node 18 to another in a peer to peer or
30 in a tic-tack fashion such that each node 18 is queried for the record corresponding to request 12. Additionally, node 18 may communicate with every other

node 18 coupled to network 14 such that a response to request 12 may be generated based on the record retrieved.

In one embodiment of the present invention, nodes 18
5 each include an index 30 that comprises a list of
information associated with the location of all nodes 18,
database units 16, and records within network 14.
Additionally, the index may include location information
associated with records within each of database units 16.
10 This index implementation would provide for the efficient
routing and retrieval of a record being queried by
request 12. Index 30 may alternatively include a list
of potential operations that may or may not be performed
by various nodes 18 or database units 16. This may allow
15 request 12 to be directed to a proper node 18 for
suitable processing before generating an appropriate
response to request 12.

In another embodiment of the present invention, one
or more nodes 18 may be operable to communicate with a
20 central server 24. Central server 24 may be coupled to
nodes 18, database units 16, or network 14. Central
server 24 may comprise a table, chart, index, or listing
that includes information associated with the location or
distribution of one or more records within one or more
25 database units 16. Request 12 may be routed directly to
central server 24 or alternatively directed to network 14
and through nodes 18 before reaching central server 24.
Alternatively, request 12 may be communicated to a web
server 26 (described below) or communicated to any other
30 suitable component in any appropriate manner according to
particular needs. Central server 24 may receive request
12 and identify where the appropriate record is found in

one or more target database units 16. Central server 24 may then directly access the record queried for or alternatively pass request 12 onto the proper node 18 associated with a target database unit 16 for further
5 processing. Alternatively, central server 24 may be eliminated entirely, whereby request 12 is directly routed to any one of nodes 18.

In yet another embodiment, each node 18 includes a unique address to be used in conjunction with a hashing
10 algorithm. The hashing algorithm may be included anywhere in data management system 10 (for example in index 30, in place thereof, or alternatively included in each node 18) to be utilized in directing request 12. Given the information contained in a record being sought,
15 the hashing algorithm may compute or otherwise determine the address of the node(s) that store the record. Accordingly, one or more nodes 18 or one or more web servers 26 may implement a hashing algorithm in order to direct request 12 to proper locations in order to
20 retrieve a record that may satisfy request 12.

In one embodiment, one or more nodes 18 are coupled to web server 26 via a link which may be any wireline, wireless, or other link suitable operable to support data communications between web server 26 and node 18 during
25 operation of data management system 10. Although a link is shown as generally coupling web server 26 to network 14, the present insertion contemplates that network 14 may communicate directly with one or more corresponding database units 16 where one or more web servers 26 are
30 included within one or more database units 16. One or more nodes 18, as well as one or more database units 16, may be integral to or separate from one or more web

5 servers 26, may operate on one or more computers, and may store any information suitable to support the operation of data management system 10 in facilitating a response to request 12. Additionally, a series of one or more web servers 26 may be coupled to (indirectly or directly) database units 16, resulting in a web server farm architecture. In this scenario, request 12 may be received by one or more web servers 26 from network 14. Nodes 18 may communicate with web servers 26 such that 10 target database units 16 may be accessed in order to obtain or otherwise retrieve a record queried for by request 12.

15 Terminal 20 is a computer operable to generate request 12 to be communicated to network 14. As used in this document, the term "computer" is intended to encompass a personal computer, workstation, network computer, wireless data port, wireless telephone, personal digital assistant, one or more microprocessors within these or other devices, a computer interface, a 20 computer program, or any other suitable processing device that is operable to generate request 12. Terminal 20 may be coupled to network 14 using one or more local area networks (LANs), metropolitan area networks (MANs), wide area networks (WANs), and/or any other appropriate 25 wireline, wireless, or other links.

Terminal 20 may include a suitable input device, such as a keypad, mouse, touch screen, microphone, or other device operable to input information. An output device coupled to terminal 20 may convey information 30 associated with the operation of data management system 10, including digital or analog data, visual information, or audio information. Terminal 20 may additionally

include one or more processors, programs, and associated memory to execute instructions and manipulate information or records within database units 16 according to the operation of data management system 10. Although only a
5 single terminal 20 is shown, a host of terminals may be coupled to network 14 that seek to retrieve a given record stored on one or more database units 16. These terminals 20 may each operate on separate computers or operate on one or more shared computers. Each of the one
10 or more terminals 20 may be a work station, personal computer (PC), network computer, personal digital assistant (PDA), wireless data port, or any other suitable computing device.

FIGURE 2 is a flowchart illustrating a series of
15 steps associated with processing request 12 with multiple database units 16 coupled to network 14. The method begins at step 100 where terminal 20 generates request 12 and communicates it to network 14. At step 102 request 12 is communicated by network 14 to any one of nodes 18
20 or an associated web server 26. The node 18, or optionally web server 26, that receives request 12 may then communicate request 12 (illustrated at step 104) to other nodes 18 or other web servers 26 in a random fashion or specifically to targeted nodes based on, for
25 example, information contained in index 30. Alternatively, node 18 or web server 26 receiving request 12 may utilize a hashing algorithm to determine where a specific piece of information is stored in any one of database units 16. In addition, node 18 or web server 26
30 receiving request 12 may query central server 24 for the location of information or a record sought by request 12. In identifying the requisite information, any one of a

number of components or elements within data management system 10 may operate to determine time-stamp information or any other data characteristic indicating or distinguishing information of records that are stored in
5 any one of database units 16. In this manner, the most recent updated copy of a record may be determined and located. This identification of a proper location or target database unit 16 and storing the record sought to be retrieved or otherwise accessed is illustrated
10 generally at step 106.

The relevant record or data stored in a targeted database unit 16 may then be accessed or retrieved by central server 24, any one or more of nodes 18, web server 26, or network 14. This access or retrieval step
15 is illustrated generally at step 108. At step 110, request 12 may then be processed and may include a simple communication of a record potentially flowing back through network 14 and to terminal 20. Alternatively, the processing of request 12 may include web server 26,
20 central server 24, or any one of nodes 18 performing some operation on request 12 in order to generate a suitable response. These operations are described above with reference to FIGURE 1.

Once request 12 is processed, any one of the
25 components or elements within data management system 10 may then return the record to one or more of database units 16. This information may be updated to reflect the operations performed or alternatively this information may be returned in its original format or structure to
30 any database unit 16. At step 112, a response to request 12 is generated and returned to network 14. The response is based on the information accessed by data management

system 10 stored in database units 16. Network 14 may then return the response to terminal 20 or alternatively further process the response before returning a result to terminal 20.

5 The present invention may be used in a host of communications environments that facilitate data storage. For example, data management system 10 may be used in wireless communication systems, asynchronous transfer mode (ATM) applications, firewall applications, e-
10 commerce environments, or any other field or area where data storage may be implemented. Additionally, although the present invention has been described with reference to web server 26, central server 24, and terminal 20, a myriad of potential devices, interfaces, hardware,
15 software, components, structure, and elements may be used in conjunction with the present invention to facilitate data storage. These objects may be coupled to or contained within network 14, or alternatively coupled to any other component within data management system 10,
20 such as database units 16 or nodes 18, for example.

In addition, the present invention may be configured such that a database management system is provided at each of nodes 18, i.e., it may be a distributed database management system. Numerous other changes,
25 substitutions, variations, alterations, and modifications may be suggested to one skilled in the art and it is intended that the present invention encompass all such changes, substitutions, variations, alterations, and modifications as falling within the spirit and scope of
30 the appended claims.